

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

PATENT APPLICATION
ATTORNEY DOCKET NO. 200310236-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Ludmila Cherkasova

Confirmation No.: 7028

Application No.: 10/619,737

Examiner: Y. Dalencourt

Filing Date: July 15, 2003

Group Art Unit: 2157

Title: SYSTEM AND METHOD HAVING IMPROVED EFFICIENCY AND RELIABILITY FOR DISTRIBUTING A FILE
AMONG A PLURALITY OF RECIPIENTS

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Transmitted herewith is the Appeal Brief in this application with respect to the Notice of Appeal filed on 6/20/2007.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) \$500.00

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

(a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d)) for the total number of months checked below:

1st Month
\$120

2nd Month
\$450

3rd Month
\$1020

4th Month
\$1590

The extension fee has already been filed in this application.

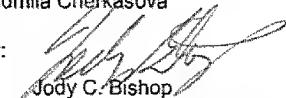
(b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account 08-2025 the sum of \$ 500. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees.

Respectfully submitted,

Ludmila Cherkasova

By:



Jody C. Bishop

Attorney/Agent for Applicant(s)

I hereby certify that this document is being transmitted to the Patent and Trademark Office via electronic filing.

Date of Transmission: June 20, 2007

Reg No. : 44,034

Date : June 20, 2007

Telephone : (214) 855-8007

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Docket No.: 200310236-1
(PATENT)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:
Ludmila Cherkasova

Application No.: 10/619,737

Confirmation No.: 7028

Filed: July 15, 2003

Art Unit: 2157

For: SYSTEM AND METHOD HAVING
IMPROVED EFFICIENCY AND
RELIABILITY FOR DISTRIBUTING A FILE
AMONG A PLURALITY OF RECIPIENTS

Examiner: Y. Dalencourt

APPEAL BRIEF

MS Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

As required under § 41.37(a), this brief is filed concurrently with the Notice of Appeal, and is in furtherance of said Notice of Appeal.

The fees required under § 41.20(b)(2) are dealt with in the accompanying
TRANSMITTAL OF APPEAL BRIEF.

This brief contains items under the following headings as required by 37 C.F.R. § 41.37 and M.P.E.P. § 1206:

- I. Real Party In Interest
- II. Related Appeals and Interferences
- III. Status of Claims
- IV. Status of Amendments

V.	Summary of Claimed Subject Matter
VI.	Grounds of Rejection to be Reviewed on Appeal
VII.	Argument
VIII.	Claims Appendix
IX.	Evidence Appendix
X.	Related Proceedings Appendix

I. REAL PARTY IN INTEREST

The real party in interest for this appeal is:

Hewlett-Packard Development Company, L.P., a Limited Partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249, Houston, TX 77070, U.S.A. (hereinafter “HPDC”). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. RELATED APPEALS, INTERFERENCES, AND JUDICIAL PROCEEDINGS

Appellant respectfully notes that the following copending applications are on appeal before the Board, which contain at least some issues that are similar to issues of the present application, which may be affected or have a bearing on the Board’s decision in this appeal: 1) Application No. 10/619,805 (hereinafter “the ‘805 application”); 2) Application No. 10/345,716 (hereinafter “the ‘716 application”); 3) Application No. 10/345,587 (hereinafter “the ‘587 application”); 4) Application No. 10/345,718 (hereinafter “the ‘718 application”); and 5) Application No. 10/345,719 (hereinafter “the ‘719 application”).

There are no other appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board’s decision in this appeal.

III. STATUS OF CLAIMS

A. Total Number of Claims in Application

There are 28 claims pending in application.

B. Current Status of Claims

1. Claims canceled: None
2. Claims withdrawn from consideration but not canceled: None
3. Claims pending: 1-28
4. Claims allowed: None
5. Claims rejected: 1-28

C. Claims On Appeal

The claims on appeal are claims 1-28

IV. STATUS OF AMENDMENTS

A Final Office Action rejecting the claims of the present application was mailed March 21, 2007. In response, Applicant did not file an Amendment After Final Rejection, but instead filed a Notice of Appeal, which this brief supports. Accordingly, the claims on appeal are those as rejected in the Final Office Action of March 21, 2007. A complete listing of the claims is provided in the Claims Appendix hereto.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the separately argued claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R. § 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and

drawings where applicable. It should be noted that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element.

According to one claimed embodiment, such as that of independent claim 1, a method of distributing a file (e.g., file F of FIG. 1) from a first node (e.g., node N₀ of FIG. 1) to a plurality of recipient nodes (e.g., recipient nodes N₁-N_k of FIG. 1) is provided. The method comprises partitioning (e.g., operational block 802 of FIG. 8) a file F into a plurality of subfiles (e.g., subfiles F₁-F_k of FIG. 1). The method further comprises performing distribution of said file F to a plurality of recipient nodes using a distribution technique that comprises

(a) attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of said first group but not all of said plurality of subfiles to any recipient node of said first group (e.g., example of FIG. 1, operational block 803 of FIG. 8, and see paragraphs 0008, 0025 and 0028), and

(b) said plurality of recipient nodes of said first group attempting to exchange their respective subfiles received from said first node, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile (e.g., example of FIGS. 2-3 and operational block 804 of FIG. 8, and see paragraphs 0008, 0026, and 0029).

The method further comprises detecting a failed node of said plurality of recipient nodes, and said distribution technique adapting to distribute all of the subfiles of said file F to each non-failed node of said plurality of recipient nodes (e.g., example of failed node N₃¹ of FIGS. 9-10C and failed node N₇¹ of FIGS. 11A-11B, operational block 1201 of FIG. 12A, and see e.g., paragraphs 0008 0030, and 0071-0083).

In certain embodiments, such as that of dependent claim 3, said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises attempting

to distribute a different subfile from said first node to each of said recipient nodes of said first group (e.g., example of FIG. 1, operational block 803A of FIG. 8, and *see* paragraph 0041).

In certain embodiments, such as that of dependent claim 5, said plurality of recipient nodes of said first group attempting to exchange their respective subfiles further comprises each of said plurality of recipient nodes attempting to establishing concurrent communication connections to every other recipient node of said first group (*see e.g.*, paragraphs 0029 and 0042-0045).

In certain embodiments, such as that of dependent claim 8, said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises said first node attempting to establish concurrent communication connections to the recipient nodes of said first group. Additionally, said distribution technique adapting comprises: responsive to said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node, said first node triggering a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate the particular subfile to said non-failed nodes (*see e.g.*, mirror node \hat{N}_0 of FIGS. 9-10C, *see e.g.*, operational blocks 1205-1208 of FIG. 12A, and *see* paragraph 0081).

According to another claimed embodiment, such as that of independent claim 17, a system comprises an origin node (e.g., node N_0 of FIG. 1) operable to partition (e.g., partitioning operation 802 of FIG. 8) a file F (e.g., file F of FIG. 1) into a plurality of subfiles (e.g., subfiles F_1-F_k of FIG. 1), wherein said plurality of subfiles correspond in number to a number of recipient nodes (e.g., recipient nodes N_1-N_k of FIG. 1) in a first group to which said file is to be distributed (*see* example of FIG. 1, and *see* paragraphs 0009 and 0039). The origin node is operable to attempt to distribute all of said plurality of subfiles to said recipient nodes, wherein said origin node attempts to distribute a different one of said plurality of subfiles to each of said recipient nodes (e.g., example of FIG. 1, operational block 803A of FIG. 8, and *see* paragraph 0041). The recipient nodes are operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile

that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node (e.g., example of FIGS. 2-3 and operational block 804 of FIG. 8, and *see* paragraphs 0008, 0026, and 0029). The origin node is also operable to detect a failed node in said first group, and the origin node is operable to manage distribution of said file *F* upon detecting a failed node in said first group in a manner such that every non-failed node of said first group receives said file *F* (e.g., example of failed node N_3^1 of FIGS. 9-10C and failed node N_i^j of FIGS. 11A-11B, operational block 1201 of FIG. 12A, and *see e.g.*, paragraphs 0008 0030, and 0071-0083).

In certain embodiments, such as that of dependent claim 20, the origin node is operable to trigger a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate a subfile to said non-failed nodes (*see e.g.*, mirror node \hat{N}_0 of FIGS. 9-10C, *see e.g.*, operational blocks 1205-1208 of FIG. 12A, and *see* paragraph 0081).

According to another claimed embodiment, such as that of independent claim 21, a method of distributing a file (e.g., file *F* of FIG. 1) from a first node (e.g., node N_0 of FIG. 1) to a plurality of recipient nodes (e.g., recipient nodes N_1-N_k of FIG. 1) is provided. The method comprises attempting to distribute a plurality of subfiles (e.g., subfiles F_1-F_k of FIG. 1) that comprise a file *F* from a first node to a first group comprising a plurality of recipient nodes (e.g., recipient nodes N_1-N_k of FIG. 1), wherein the first node attempts to distribute at least one subfile to each recipient node of said first group but not all of said plurality of subfiles are distributed from the first node to any of the recipient nodes of said first group (e.g., example of FIG. 1, operational block 803 of FIG. 8, and *see* paragraphs 0008, 0025 and 0028). The plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile (e.g., example of FIGS. 2-3 and operational block 804 of FIG. 8, and *see* paragraphs 0008, 0026, and 0029). The method further comprises detecting whether one of said plurality of recipient nodes of said first group has failed, and if a recipient node of said first group has failed, managing the distribution

of the plurality of subfiles to detour their distribution around the failed node such that the file F is distributed to each non-failed node of said plurality of recipient nodes (e.g., example of failed node N_3^1 of FIGS. 9-10C and failed node N_i^j of FIGS. I1A-11B, operational block 1201 of FIG. 12A, and *see e.g.*, paragraphs 0008 0030, and 0071-0083).

In certain embodiments, such as that of dependent claim 22, said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group (e.g., example of FIG. 1, operational block 803A of FIG. 8, and *see* paragraph 0041).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-28 are provisionally rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1-47 of copending Application 10/345,587 (hereinafter “the ‘587 application”) in view of U.S. Patent Publication 2004/0088380 to Chung et al. (hereinafter “*Chung*”).

Claims 1-28 are rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent 6,970,939 to Siew Yong Sim (hereinafter “*Sim*”) in view of U.S. Patent 6,477,583 to Zayas et al. (hereinafter “*Zayas*”), and further in view of *Chung*.

VII. ARGUMENT

Appellant respectfully traverses the outstanding rejections of the pending claims, and requests that the Board reverse the outstanding rejections in light of the remarks contained herein. The claims do not stand or fall together. Instead, Appellant presents separate arguments for various independent and dependent claims. Each of these arguments is separately argued below and presented with separate headings and sub-heading as required by 37 C.F.R. § 41.37(c)(1)(vii).

A. Rejections Under Obviousness-type Double Patenting

Claims 1-28 are provisionally rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1-47 of the '587 application in view of *Chung*. Appellant respectfully traverses this rejection for the reasons stated below.

An obviousness-type double patenting rejection should make clear the differences between the inventions defined by the conflicting claims—a claim in the patent compared to a claim in the application, and the reasons why a person of ordinary skill in the art would conclude that the invention defined in the claim at issue would have been an obvious variation of the invention defined in a claim in the patent. M.P.E.P. § 804(II)(B)(1). Thus, to establish a proper grounds of obviousness-type double patenting, the Examiner must apply the *Graham* obviousness analysis. The Final Office Action fails to properly establish such obviousness-type double patenting for the reasons discussed below.

The Final Office Action concedes at page 6 thereof that claims 1, 22, and 36 of the '587 application fails to recite that at least one recipient node of a first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of the first group before the at least one recipient node fully receives its respective subfile. However, the Final Office Action contends that *Chung* teaches this element, *see* page 7 of the Final Office Action. Appellant respectfully disagrees. For the reasons discussed further herein below, *Chung* does not teach or suggest this element, and thus the Final Office Action fails to establish a proper obviousness-type double patenting rejection.

Further, the Final Office Action fails to establish any reasoning why one of ordinary skill in the art would have been motivated to combine the teaching of *Chung* with the claims of the ‘587 application. To properly establish obviousness, sufficient reasons that would have prompted one of ordinary skill in the art to combine/modify the prior art teachings in a manner that would arrive at the claimed invention must be present. For instance, it is “important to identify a reason that would have prompted a person of ordinary skill … to have combined the [prior art] elements in the way the claimed invention does”. *KSR Int'l Co. v. Teleflex, Inc.*, No 05-1350, 550 U.S. _____ (2007). Thus, sufficient reasoning for one of ordinary skill in the art to have been motivated to combine *Chung* with the ‘587 application must be established. The Final Office Action fails to establish such reasoning.

Further, the Final Office Action asserts on page 7 thereof that claims 2-16, 18-20, and 22-28 of the instant application are the same as claims of “copending Application No. 10/619,737”. Appellant respectfully notes that the present application is Application No. 10/619,737, and thus this statement by the Examiner fails to provide any support whatsoever for an obvious-type double patenting rejection.

Accordingly, Appellant respectfully requests that the obviousness-type double patenting rejection be overturned.

B. Rejections Under 35 U.S.C. §103 over *Sim* in view of *Zayas* and *Chung*

Claims 1-28 are rejected under 35 U.S.C. §103(a) as being unpatentable over *Sim* in view of *Zayas*, and further in view of *Chung*. Appellant respectfully traverses these rejections below.

The test for non-obvious subject matter is whether the differences between the subject matter and the prior art are such that the claimed subject matter as a whole would have been obvious to a person having ordinary skill in the art to which the subject matter pertains. The United States Supreme Court in Graham v. John Deere and Co., 383 U.S. 1 (1966) set forth the factual inquiries which must be considered in applying the statutory test: (1) determining of the scope and content of the prior art; (2) ascertaining the differences between the prior art and the claims at issue; and (3) resolving the level of ordinary skill in the pertinent art. As discussed further hereafter, Appellant respectfully asserts that the claims include non-obvious differences over the cited art.

As discussed further below, the rejections should be overturned because when considering the scope and content of the applied *Sim*, *Zayas*, and *Chung* references there are significant differences between the applied combination and claims 1-28, as the applied combination fails to disclose all elements of these claims. Thus, considering the lack of disclosure in the applied combination of all elements of claims 1-28, one of ordinary skill in the art would not find these claims obvious under 35 U.S.C. §103, and therefore the rejections should be overturned. Additionally, insufficient reasoning exists for one of ordinary skill in the art to combine the *Sim*, *Zayas*, and *Chung* teachings in a manner as to arrive at the elements as recited by claims 1-28, and the rejections should therefore be overturned for this further reason, as also discussed below.

1. The Applied Combination Fails to Teach or Suggest All Claim Elements**Independent Claim 1 and Dependent Claims 2, 4, 6-7, and 9-16**

Independent claim 1 recites:

A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

partitioning a file *F* into a plurality of subfiles;

performing distribution of said file *F* to a plurality of recipient nodes using a distribution technique that comprises

(a) attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of said first group but not all of said plurality of subfiles to any recipient node of said first group, and

(b) said plurality of recipient nodes of said first group attempting to exchange their respective subfiles received from said first node, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;

detecting a failed node of said plurality of recipient nodes; and

said distribution technique adapting to distribute all of the subfiles of said file *F* to each non-failed node of said plurality of recipient nodes. (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest at least the above-emphasized element of claim 1. As discussed below, the applied combination fails to teach or suggest at least: 1) a plurality of recipient nodes attempting to exchange their respective subfiles received from a first node; and 2) at least one recipient node beginning to communicate a portion of its respective subfile that it is receiving from the first node to at least one other recipient node before the at least one recipient node fully receives its respective subfile. The disclosure of each reference is addressed below to illustrate that the combination of the disclosures fails to teach or suggest at least this element of claim 1.

First, *Sim* fails to teach or suggest the above-emphasized elements of claim 1. In *Sim*, a large payload file is partitioned into a plurality of subfiles. The subfiles are all stored to an originating node. Certain ones of the subfiles are then distributed to other nodes for serving to

end-user clients. However, *Sim* does not teach or suggest a distribution technique in which a first node attempts to distribute the plurality of subfiles that make up a file *F* from the first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of the first group but not all of the plurality of subfiles to any recipient node of the first group. Further, *Sim* fails to teach or suggest that the plurality of recipient nodes of the first group attempt to exchange their respective subfiles received from the first node. Rather, in *Sim* the subfiles desired by one recipient node are obtained from a nearest node possessing the desired subfiles. As discussed further below, no exchange of respective subfiles received from a first node is attempted between the recipient nodes of a first group in *Sim*.

Sim explains its system for distributing a large payload file at column 14, lines 10-42 as follows:

A content provider uploads a large payload file to a single content management server using content publishing and management tools running on a content provider client system. After receiving the file, the CMS processes the file and breaks it down, if required, into track files (a.k.a. linear files). A linear file comprises a file that maintains the order associated with the substance (i.e., substantive content) of the file. If, for example, the linear file contained a movie, the beginning of that file would include the beginning portions of the movie. Similarly, the middle and end portions of the movie would be located at the middle and end of the linear file. Linear files are desired because it is easier to reassemble such files using linear superposition, for example. Some media files are non-linear, that is, they contain multiple tracks such that the first part of the movie, for example, is not stored in the beginning of the file. After breaking the file down to linear (i.e., track) files, the CMS transfers the file to the distribution server it is connected to. The distribution server further breaks the track files down to block files, as desired for storage. The block files may subsequently be stored in local storage locations 711-713, for example. A file distribution protocol (e.g., FDP) command is subsequently used to distribute (i.e., replicate) the file, or selected portions thereof, to other distribution server nodes within the scalable content delivery network. For initial replication, the entire block files need not be stored in all nodes however a master copy may be maintained completely in one node (typically the originating node). The FDP includes commands to facilitate file transfers and manipulations within the SCDN. The size of the blocks affects the performance of both content distribution and content delivery and is discussed later in this document.

Sim illustrates its distribution technique in FIG. 13, which it explains at col. 20, line 34 – col. 21, line 24 as follows:

FIG. 13 is an illustrative embodiment of the distribution of a large payload file within an SCDN. A content provider uploads a large payload file into the content management server (CMS) 570, which is connected to node B of the SCDN, using any content publishing and management software running on the content provider's client system (CPC) 530. The content provider also uploads the distribution criteria onto CMS 570. Content management server 570, as previously described, divides the uploaded file into track files and issues a command similar to the FDP "put" command for each track file to the distribution server located in node B. In other embodiments, the CMS may be connected to any node of the SCDN. At node B, the DS divides the track files into block files for local storage. The full copy of the file is shown at Node B as a filled in dot. The CMS then issues an FDP command of the type "distribute" to the distribution server at node B. In response to the distribute command, the DS issues a command to its neighboring nodes A, D, and E to replicate the content (e.g., using the "replicate" command of the FDP). Node D examines the replicate packet and decides its not supposed to have the content thus it passes the replicate command to its neighbor, node H. Nodes A, E, and H examine the replicate packet and decide they all match the distribution criteria (i.e., they are "qualified" nodes). When ready, nodes A, E, and H issue commands to retrieve a portion of the file from the nearest node (e.g., node B) in the SCDN. Nodes E and H are leaf nodes thus they do not propagate the replicate command. However, node A is the root node with child nodes B and C. Node A may not send the replicate command back to node B, because it is the originating node. However, node A may send the replicate request to node C. Node C checks the distribution criteria and decides it's a qualified node therefore it retrieves a portion of the file from the nearest nodes (e.g., the nearest of nodes A, B, E, and H) containing the needed data. Node C subsequently sends the replicate command to nodes F and G. Node F is qualified thus it retrieves a portion of the file from the nearest nodes having the data (e.g. nodes B or C). Nodes G and H are not qualified thus they receive nothing. Node G is a terminating node because the rolled-up attribute of its branch does not satisfy the distribution criteria. This initial replication process continues until all the qualified nodes in SCDN are at least partially populated. In one or more embodiments, the same portion (e.g., blocks) of the large payload file is contained in at least one node of the SCDN. Preferably, a plurality of nodes maintains the same portion thereby creating redundancy and preventing loss of any portion of the large payload file when one or more nodes or storage volumes become unavailable. For example, when a storage volume (or device) becomes unavailable (i.e., lost), a DS at that station need not take any special action to recover contents of the damaged volume since the portions of large payload files stored and hence lost in that volume are automatically downloaded from other network nodes upon demand to service a user request. The distribution servers

also relay control information of a failed station to neighbors of the failed station to prevent improper termination of control commands.

Sim does not teach that the recipient nodes of a first group exchange their respective subfiles. Rather, as shown in FIG. 13, node B (an origin node) contains all of the subfiles. Node B sends a “distribute” command to its neighboring nodes A, D, and E, which is propagated through the other nodes H, C, and F-I, and such nodes determine based on a distribution criteria whether they are to obtain any of the subfiles. If they are to obtain any of the subfiles, the nodes A and C-I retrieve those subfiles that they are to obtain from a nearest node possessing such subfiles.

As shown in FIG. 13, nodes A, C, E, F, and H each receive some of the subfiles. However, *Sim* does not teach that any of the recipient nodes A, C, E, F, and H exchange their respective subfiles. While a given recipient node may receive its desired subfiles from another of the recipient nodes that is nearest it, *Sim* does not teach that the recipient nodes “exchange” their respective subfiles.

Further, in the example of FIG. 13 of *Sim*, assuming that nodes A, E, and H may be considered a “first group” which receive subfiles from origin node B, *Sim* provides no teaching whatsoever that the recipient nodes A, E, and H exchange their respective subfiles received from origin node B. Rather, in *Sim*, whatever subfiles are desired by recipient nodes A, E, and H, those recipient nodes obtain from a nearest node possessing the desired subfiles (e.g., origin node B). That is, *Sim* provides no teaching that nodes A, E, and H each receive subfiles from origin node B, and then attempt to exchange their respective subfiles received from node B. For example, in FIG. 13, it appears that nodes A and E each receive subfiles from node B. *Sim* provides no teaching that nodes A and E then exchange their respective subfiles received from node B with each other. Even in the event that node E later desires a subfile that is possessed by node A (and assuming that node A is nearest node E), node A may send such desired subfile to node E, but *Sim* provides no teaching of an “exchange” in which node E also sends a subfile to node A.

Thus, *Sim* does not teach or suggest at least this “exchanging” element of claim 1. The Examiner appears to be ignoring the express recitation of “said plurality of recipient nodes of said first group attempting to exchange their respective subfiles received from said first node” (emphasis added) of claim 1 in maintaining that *Sim* discloses this element. In response to the above arguments, the Final Office Action asserts on page 3 thereof:

Sim discloses that Nodes A, E, and H examine the replicate packet and decide they all match the distribution criteria (i.e., they are “qualified” nodes). When ready, nodes A, E, and H issue commands to retrieve a portion of the files from the nearest node (i.e., node B) in the SCDN. Nodes E and H are leaf nodes thus they do not propagate the replicate command back to node B, because it is the originating node. However, node A may send the replicate request to node C, thus, Sim does teach in a broader sense the claimed limitation of “of nodes A and E exchanging their respective subfiles received from node B (see col. 14, lines 25-39; col. 16, lines 7-16; col. 20, line 34 through col. 21, line 24).

The above explanation by the Office Action in no way supports a conclusion that *Sim* discloses a plurality of recipient nodes of a first group attempting to exchange their respective subfiles received from a first node, as recited by claim 1. Merely because node A “may send the replicate request to node C” in no way means that nodes A and E exchange their respective subfiles received from node B, as asserted by the Examiner. Further, even if node A sends a replicate request on to node C and if node C determines that it matches a distribution criteria (is a qualified node), *Sim* does not teach that nodes A and C exchange their respective subfiles that they received from a first node (e.g., node B); but instead, node C would request any desired subfiles from a nearest node (which may or may not be node A) and certainly does not, under *Sim*’s teaching, “exchange” respective subfiles with another node.

For instance, as discussed above, assuming that nodes A, E, and H may be considered a “first group” which receive subfiles from origin node B, *Sim* provides no teaching whatsoever that the recipient nodes A, E, and H exchange their respective subfiles received from origin node B. Rather, in *Sim*, whatever subfiles are desired by recipient nodes A, E, and H, those recipient nodes obtain from a nearest node possessing the desired subfiles (e.g., origin node B), as the Examiner appears to concede. That is, *Sim* provides no teaching that nodes A, E, and H each receive subfiles from origin node B, and then attempt to exchange their respective subfiles

received from node B. For example, in FIG. 13, it appears that nodes A and E each receive subfiles from node B. *Sim* provides no teaching that nodes A and E then exchange their respective subfiles received from node B with each other. Even in the event that a node, such as node C, later desires a subfile that is possessed by node A (and assuming that node A is nearest node C), node A may send such desired subfile to node C, but *Sim* provides no teaching of an “exchange” in which node C also sends a subfile to node A.

Indeed, *Sim* teaches away from the recipient nodes exchanging their respective subfiles received from the origin node because *Sim* expressly attempts to avoid each recipient node receiving all of the subfiles. For instance, *Sim* teaches at col. 8, lines 59-62 that “[e]ach node at the edge of the network embodying aspects of the invention is configured to appear as if it has the large file stored locally when portions of the file are really stored on other nodes located throughout the network.” At col. 9, lines 2-4, *Sim* further explains that the “end result is that each network node has access to numerous large data files without having to store each of those data files locally.” As shown in the example of FIG. 13 of *Sim*, the origin node B contains all of the subfiles, while the recipient nodes A, C, E, F, and H each receive some of the subfiles, but not all of the subfiles. *Sim* does not provide any teaching or suggestion of the nodes exchanging their respective subfiles received from the first node.

While the Final Office Action first appears to assert on page 3 thereof that *Sim* discloses the nodes exchanging their respective subfiles received from the first node (i.e., “*Sim* does teach in a broader sense the claimed limitation of nodes A and E exchanging their respective subfiles received from node B”), the Final Office Action goes on to concede that *Sim* indeed fails to teach or suggest the plurality of recipient nodes of the first group attempting to exchange their respective subfiles received from the first node, *see* pages 3 and 8 of the Final Office Action. However, the Final Office Action then asserts that *Zayas* teaches this element of the claim. Appellant disagrees and maintains that *Zayas* also fails to disclose this element as discussed below.

Zayas is directed to an infrastructure in which volumes are replicated on each of a plurality of servers. The infrastructure permits different “replication modules” to be utilized for

managing the replication (e.g., distributing updates to the volumes, etc.) of different volumes. For instance, FIG. 2A of *Zayas* shows an example in which a replication module A 265A is utilized for managing replication of a volume V3, while a different replication module B 265B is utilized for managing replication of volumes V1 and V2. Thus, *Zayas* explains that if one of the replication modules employed incurs large overhead for keeping the volumes that it manages consistent, this does not lead necessarily to a large overhead for the volumes managed by a different replication module, *see e.g.* column 4, lines 26-35 of *Zayas*.

Accordingly, *Zayas* describes a system in which entire volumes of files are replicated onto a plurality of different servers. *Zayas* is not concerned, however, with how the files are distributed to the different servers. For instance, presumably one server in *Zayas* may simply distribute a full file to each of the other servers to which the file is to be replicated. *Zayas* is instead concerned with a system in which different modules can be used for managing updates to different volumes in order to maintain consistency in the volumes across the different servers to which they are stored. Accordingly, *Zayas* also fails to teach or suggest a plurality of recipient nodes attempting to exchange their respective subfiles received from a first node. Thus, the Final Office Action has failed to establish that the combination of references teaches or suggests at least a plurality of recipient nodes of a first group attempting to exchange their respective subfiles received from a first node, as recited by claim 1.

Additionally, the Final Office Action concedes on page 9 thereof that neither *Sim* nor *Zayas* teaches or suggests at least one recipient node beginning to communicate a portion of its respective subfile that it is receiving from the first node to at least one other recipient node before the at least one recipient node fully receives its respective subfile. However, the Office Action contends that *Chung* provides this element, citing to paragraphs 0008-0009 and 0016 of *Chung*, *see* item 13 on page 9 of the Final Office Action. Appellant disagrees for the reasons discussed below.

While *Chung* mentions that subfiles may be distributed to different servers, and may then be streamed in parallel from the servers to an end user, *Chung* makes no mention whatsoever of a given recipient node of a subfile in a group beginning to communicate a portion of the subfile

that the given recipient node is receiving to another node in the group before the given recipient node fully receives the subfile. First, the servers to which the subfiles are distributed do not exchange their respective subfiles, and *Chung* provides no teaching whatsoever of any of the servers beginning to communicate a subfile to another server before the subfile is fully received. Further, while subfiles may be streamed to an end user from the servers, *Chung* does not teach or suggest that a server begins streaming a subfile to an end user before the subfile is fully received by the server. Additionally, *Chung* provides no teaching that the end user communicates a received subfile to another recipient node, and certainly fails to teach or suggest that the end user begins communicating the received subfile to another recipient node before the end user fully receives the subfile.

The Examiner appears to contend that the servers of *Chung* provide the recited first group comprising a plurality of recipient nodes. However, *Chung* provides no teaching of one server communicating a subfile that it receives to another of the servers. Further, if the Examiner further considers an end user of *Chung* as a recipient node in the recited first group, *Chung* likewise fails to teach or suggest that the end user communicates the subfile that it receives to another node. Thus, irrespective of how a given end user may receive a subfile from the servers (e.g., via streaming or other download technique), *Chung* provides no disclosure of the given end user communicating the subfile that it receives to another recipient nodes. *Chung* certainly provides no teaching of one recipient node (the end user or a server) beginning to communicate a portion of its respective subfile that it is receiving to another of the recipient nodes (another server or end user) before the one recipient node fully receives its respective subfile.

As mentioned above, the Final Office Action cites to paragraph 0016 of *Chung* as teaching this element of claim 1. Paragraph 0016 of *Chung* merely provides:

One aspect of the invention involves dividing a single file into multiple files or sub-files. A sub-file has a file name and other file attributes, and is treated by the operating system's file system as just another file. The divided files or sub-files may then be distributed and stored onto one or more servers. When an end user wants the file to be delivered in a streaming fashion, the sub-files can be transmitted in parallel and simultaneously from one or more servers, which increases the rate at which data can be delivered.

This portion of *Chung* mentions that the subfiles can be distributed (from an origination node that contains the full file) to the various servers, such as servers A-D of FIGURE 4. This portion of *Chung* further mentions that the subfiles from the servers can be distributed to an end user (client) in parallel and delivered to the end user in a streaming fashion. Thus, if a client such as client 610 in FIGURE 6 of *Chung* requests a file, the various subfiles S1-S4 may be sent to the client in parallel from the different servers A-D.

As discussed above, *Chung* provides no teaching or suggestion of a recipient node (either a server or an end user) communicating its subfile to another recipient node. Further, *Chung* provides no teaching or suggestion of one recipient node (either a server or an end user) beginning to communicate a portion of its respective subfile to another of the recipient nodes before the one recipient node fully receives its respective subfile. For instance, even when servers A-D communicate their respective subfiles to the client node 610, *Chung* makes no mention of the servers beginning such communication of a subfile from the server to the client before the subfile is fully received by the server.

Thus, *Chung* fails to teach or suggest the above element of claim 1. Neither *Sim* nor *Zayas* is relied upon as teaching such element, nor do they do so.

In response to the above arguments, the Examiner asserts at page 5 of the Final Office Action that “Applicants are interpreting the claims very narrow without considering the broad teaching of the reference used in the rejection.” Appellant disagrees. Appellant has not interpreted the claims very narrowly, but has instead clearly pointed out expressly recited language of claim 1 and has explained in detail how the teachings of the applied references simply provide no disclosure of such expressly recited elements.

In view of the above, the applied combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all elements of claim 1, and thus the rejection of claim 1 should be overturned.

Claims 2, 4, 6-7, and 9-16 each depend either directly or indirectly from independent claim 1, and are thus likewise believed to be allowable at least based on their dependency from

claim 1 for the reasons discussed above. Accordingly, Appellant respectfully requests that the rejection of claims 2, 4, 6-7, and 9-16 also be overturned.

Dependent Claim 3

Dependent claim 3 depends from claim 1, and thus inherits all of the limitations of claim 1 in addition to its own supplied limitations. It is respectfully submitted that dependent claim 3 is allowable at least because of its dependence from claim 1 for the reasons discussed above.

Claim 3 further recites “wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises: attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group.” (Emphasis added). The combination of *Sim*, *Zayas*, and *Chung* further fails to teach or suggest this element of claim 3. *Zayas* and *Chung* are not relied upon as teaching or suggesting this element (*see* page 10 of Final Office Action), nor do they do so. Further, *Sim* fails to teach or suggest this element, as discussed below.

While *Sim* teaches distributing subfiles from a first node, such as node B in FIG. 13, to a plurality of recipient nodes, *Sim* provides no teaching or suggestion of attempting to distribute a different subfile from the first node to each of the recipient nodes of a first group. For instance, it appears from the shading shown in FIG. 13 that node B distributes certain subfiles to recipient node A, and some of the same subfiles are also distributed from node B to recipient node E. That is, it appears that node E receives from node B a subset of the subfiles that are sent to node A. *Sim* provides no teaching of attempting to distribute a different subfile from the first node to each of the recipient nodes of a first group.

Thus, for this further reason, the rejection of claim 3 is improper and should be overturned.

Dependent Claim 5

Dependent claim 5 depends from claim 1, and thus inherits all of the limitations of claim 1 in addition to its own supplied limitations. It is respectfully submitted that dependent claim 5 is allowable at least because of its dependence from claim 1 for the reasons discussed above.

Claim 5 further recites “wherein said plurality of recipient nodes of said first group attempting to exchange their respective subfiles further comprises: each of said plurality of recipient nodes attempting to establishing concurrent communication connections to every other recipient node of said first group.” (Emphasis added). The combination of *Sim*, *Zayas*, and *Chung* further fails to teach or suggest this element of claim 5. *Zayas* and *Chung* are not relied upon as teaching or suggesting this element (see page 10 of Final Office Action), nor do they do so. Further, *Sim* fails to teach or suggest this element, as discussed below.

Sim does not teach that each of the recipient nodes establish concurrent communication connections to every other recipient node.

The Final Office Action cites (at page 10 thereof) to col. 11, lines 36-47 of *Sim* as teaching the above element of claim 5. The cited portion of *Sim* provides:

In one embodiment of the invention, the system accomplishes this by breaking the large payload file into multiple portions and storing those portions in locations (e.g., nodes) distributed throughout the network. The portions stored throughout the network are distributed utilizing a flow optimization technique that provides for the intelligent management of the large payload files. Thus, portions of the large payload file are stored in locations that minimize the amount of time it takes to deliver the portion to the end-user system. These locations minimize the latency associated with delivering the file to the end-user system and are referred to herein as the edge of the network.

As can be seen, the cited portion of *Sim* fails to teach, in any way, the above element of claim 5. The above portion of *Sim* fails to mention any nodes that exchange subfiles, and certainly fails to teach that each of a plurality of recipient nodes establish concurrent communication connections with every other one of the plurality of recipient nodes.

Accordingly, for this further reason, the rejection of claim 5 is improper and should be overturned.

Dependent Claim 8

Dependent claim 8 depends indirectly from claim 1, and thus inherits all of the limitations of claim 1 in addition to its own supplied limitations. It is respectfully submitted that dependent claim 8 is allowable at least because of its dependence from claim 1 for the reasons discussed above.

Claim 8 further recites:

wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises said first node attempting to establish concurrent communication connections to the recipient nodes of said first group, and wherein said distribution technique adapting comprises:

responsive to said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node, said first node triggering a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate the particular subfile to said non-failed nodes. (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* further fails to teach or suggest this element of claim 8. *Zayas* and *Chung* are not relied upon as teaching or suggesting this element (see page 11 of Final Office Action), nor do they do so. Further, *Sim* fails to teach or suggest this element, as discussed below.

Sim does not teach that responsive to a first node (e.g., node B in *Sim*) detecting a failed node in a first group of recipient nodes such that the first node is unable to communicate a particular subfile to such failed node, the first node triggering a mirror node to establish concurrent communication connections with non-failed nodes of the first group to communicate the particular subfile to such non-failed nodes.

The Final Office Action cites (at page 11 thereof) to col. 21, lines 1-37 of *Sim* as teaching the above element of claim 8. The cited portion of *Sim* provides:

(e.g., the nearest of nodes A, B, E, and H) containing the needed data. Node C subsequently sends the replicate command to nodes F and C. Node F is qualified thus it retrieves a portion of the file from the nearest nodes having the data (e.g. nodes B or C). Nodes G and/or are not qualified thus they receive nothing. Node G is a terminating node because the rolled-up attribute of its branch does not satisfy the distribution criteria. This initial replication process continues until all the qualified nodes in SCDN are at least partially populated. In one or more embodiments, the same portion (e.g., blocks) of the large payload file is contained in at least one node of the SCDN. Preferably, a plurality of nodes maintains the same portion thereby creating redundancy and preventing loss of any portion of the large payload file when one or more nodes or storage volumes become unavailable. For example, when a storage volume (or device) becomes unavailable (i.e., lost), a DS at that station need not take any special action to recover contents of the damaged volume since the portions of large payload files stored and hence lost in that volume are automatically downloaded from other network nodes upon demand to service a user request. The distribution servers also relay control information of a failed station to neighbors of the failed station to prevent improper termination of control commands.

During normal operation, a Distribution Server sends FDP commands, such as replicate, info, search, and clean commands that are forwarded to all or part of the network, through other Distribution Servers in the immediate neighbor stations in its control path. For example, when a Distribution Server receives an FDP command such as replicate or info, it sends the command to its neighbor DSs based on the FDP distribution criteria. In the situation when one of the neighbor stations is failed, the DS keeps the job in its job queue, and repeatedly retries until the job is successfully completed. At the same time, the DS temporarily assumes the role of the DS in the failed station by forwarding the FDP command to the neighbor DSs of the failed station.

As can be seen, the cited portion of *Sim* fails to teach, in any way, the above element of claim 8. The above portion of *Sim* mentions that when one of the neighbor stations is failed, a distribution server (DS) keeps the job in its job queue, and repeatedly retries until the job is successfully completed, and the DS temporarily assumes the role of the DS in the failed station by forwarding the FDP command to the neighbor DSs of the failed station. This does not disclose that when a first node (e.g., node B in *Sim*) detects a failed node in a first group of recipient nodes, the first node triggers a mirror node to establish concurrent communication connections with non-failed nodes of the first group to communicate the particular subfile to such non-failed nodes. The above-portion of *Sim* makes no mention of triggering a mirror node, nor does it mention establishing concurrent communication connections with non-failed nodes of

a first group to communicate a particular subfile that was unable to be communicated to a failed node.

Accordingly, for this further reason, the rejection of claim 8 is improper and should be overturned.

Independent Claim 17 and Dependent Claims 18-19

Independent claim 17 recites:

A system comprising:

an origin node operable to partition a file F into a plurality of subfiles, wherein said plurality of subfiles correspond in number to a number of recipient nodes in a first group to which said file is to be distributed;

said origin node operable to attempt to distribute all of said plurality of subfiles to said recipient nodes, wherein said origin node attempts to distribute a different one of said plurality of subfiles to each of said recipient nodes;

said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node;

said origin node operable to detect a failed node in said first group; and

said origin node operable to manage distribution of said file F upon detecting a failed node in said first group in a manner such that every non-failed node of said first group receives said file F . (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest at least the above-emphasized element of claim 17. As discussed above with claim 1, the applied combination fails to teach or suggest at least: 1) said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node, and 2) at least one recipient node that begins communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node. Further, the applied combination fails to teach or suggest such an attempt to exchange the subfiles “such that each recipient node obtains all of said plurality of subfiles”. Indeed, as discussed further herein, *Sim* and *Chung* each

expressly teach away from attempting to have each recipient node obtain all of the plurality of subfiles.

Additionally, as discussed above with dependent claim 3, the combination of *Sim*, *Zayas*, and *Chung* further fails to teach or suggest that an origin node attempts to distribute a different one of said plurality of subfiles to each of said recipient nodes. *Zayas* and *Chung* are not relied upon as teaching or suggesting this element (*see* treatment of claim 3 on page 10 of Final Office Action), nor do they do so. Further, *Sim* fails to teach or suggest this element, as discussed below.

While *Sim* teaches distributing subfiles from a first node, such as node B in FIG. 13, to a plurality of recipient nodes, *Sim* provides no teaching or suggestion of attempting to distribute a different subfile from the first node to each of the recipient nodes of a first group. For instance, it appears from the shading shown in FIG. 13 that node B distributes certain subfiles to recipient node A, and some of the same subfiles are also distributed from node B to recipient node E. That is, it appears that node E receives from node B a subset of the subfiles that are sent to node A. *Sim* provides no teaching of attempting to distribute a different subfile from the first node to each of the recipient nodes of a first group.

Thus, the applied combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all elements of claim 17, and thus the rejection of claim 17 should be overturned.

Claims 18-19 each depend either directly or indirectly from independent claim 17, and are thus likewise believed to be allowable at least based on their dependency from claim 17 for the reasons discussed above. Accordingly, Appellant respectfully requests that the rejection of claims 18-19 also be overturned.

Dependent Claim 20

Dependent claim 20 depends from claim 17, and thus inherits all of the limitations of claim 17 in addition to its own supplied limitations. It is respectfully submitted that dependent

claim 20 is allowable at least because of its dependence from claim 17 for the reasons discussed above.

Claim 20 further recites: “wherein said origin node is operable to trigger a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate a subfile to said non-failed nodes.” (Emphasis added). The combination of *Sim*, *Zayas*, and *Chung* further fails to teach or suggest this element of claim 20. *Zayas* and *Chung* are apparently not relied upon as teaching or suggesting this element (*see* treatment of claim 8 on page 11 of Final Office Action), nor do they do so. Further, *Sim* fails to teach or suggest this element, as discussed below.

Sim does not teach that an origin node (e.g., node B in *Sim*) triggers a mirror node to establish concurrent communication connections with non-failed nodes of the first group to communicate a subfile to such non-failed nodes. *Sim* mentions at col. 21, lines 1-37 thereof that when one of the neighbor stations is failed, a distribution server (DS) keeps the job in its job queue, and repeatedly retries until the job is successfully completed, and the DS temporarily assumes the role of the DS in the failed station by forwarding the FDP command to the neighbor DSs of the failed station. However, this does not disclose that an origin node (e.g., node B in *Sim*) is operable to trigger a mirror node to establish concurrent communication connections with non-failed nodes of the first group to communicate a subfile to such non-failed nodes. *Sim* makes no mention of triggering a mirror node, nor does it mention establishing concurrent communication connections with non-failed nodes of a first group to communicate a subfile thereto.

Accordingly, for this further reason, the rejection of claim 20 is improper and should be overturned.

Independent Claim 21 and Dependent Claims 23-28

Independent claim 21 recites:

A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

attempting to distribute a plurality of subfiles that comprise a file F from a first node to a first group comprising a plurality of recipient nodes, wherein the first node attempts to distribute at least one subfile to each recipient node of said first group but not all of said plurality of subfiles are distributed from the first node to any of the recipient nodes of said first group;

said plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;

detecting whether one of said plurality of recipient nodes of said first group has failed; and

if a recipient node of said first group has failed, managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that the file F is distributed to each non-failed node of said plurality of recipient nodes. (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest at least the above-emphasized element of claim 21. As discussed above with claim 1, the applied combination fails to teach or suggest at least: “said plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile”.

Thus, the applied combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all elements of claim 21, and thus the rejection of claim 21 should be overturned.

Claims 23-28 each depend either directly or indirectly from independent claim 21, and are thus likewise believed to be allowable at least based on their dependency from claim 21 for the reasons discussed above. Accordingly, Appellant respectfully requests that the rejection of claims 23-28 also be overturned.

Dependent Claim 22

Dependent claim 22 depends from claim 21, and thus inherits all of the limitations of claim 21 in addition to its own supplied limitations. It is respectfully submitted that dependent claim 22 is allowable at least because of its dependence from claim 21 for the reasons discussed above.

Claim 22 further recites “wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises: attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group.” (Emphasis added). The combination of *Sim*, *Zayas*, and *Chung* further fails to teach or suggest this element of claim 22. *Zayas* and *Chung* are not relied upon as teaching or suggesting this element (see page 10 of Final Office Action), nor do they do so. Further, *Sim* fails to teach or suggest this element, as discussed below.

While *Sim* teaches distributing subfiles from a first node, such as node B in FIG. 13, to a plurality of recipient nodes, *Sim* provides no teaching or suggestion of attempting to distribute a different subfile from the first node to each of the recipient nodes of a first group. For instance, it appears from the shading shown in FIG. 13 that node B distributes certain subfiles to recipient node A, and some of the same subfiles are also distributed from node B to recipient node E. That is, it appears that node E receives from node B a subset of the subfiles that are sent to node A. *Sim* provides no teaching of attempting to distribute a different subfile from the first node to each of the recipient nodes of a first group.

Thus, for this further reason, the rejection of claim 22 is improper and should be overturned.

2. Insufficient Reasoning for One of Ordinary Skill in the Art to Combine the References in the Manner Applied

Further, to make a proper rejection under 35 U.S.C. §103(a), sufficient reasons that would have prompted one of ordinary skill in the art to combine/modify the prior art teachings in a manner that would arrive at the claimed invention must be present. For instance, it is “important to identify a reason that would have prompted a person of ordinary skill … to have combined the [prior art] elements in the way the claimed invention does”. *KSR Int'l Co. v. Teleflex, Inc.*, No 05-1350, 550 U.S. ____ (2007).

In this case, insufficient reasoning exists for one of ordinary skill in the art to combine the teachings of *Sim*, *Zayas*, and *Chung* in the manner relied upon by the Final Office Action. The Final Office Action alleges that one of ordinary skill in the art would be motivated to combine *Sim* and *Zayas* to arrive at a system in which a plurality of nodes of a first group attempt to exchange their respective subfiles received from a first node (or origin node), *see e.g.*, the rejection of claim 1.

However, *Sim* expressly teaches away from any such system that attempts a distribution that results in each recipient node obtaining all of the subfiles, as discussed above with claim 1. For instance, *Sim* teaches at col. 8, lines 59-62 that “[e]ach node at the edge of the network embodying aspects of the invention is configured to appear as if it has the large file stored locally when portions of the file are really stored on other nodes located throughout the network.” At col. 9, lines 2-4, *Sim* further explains that the “end result is that each network node has access to numerous large data files without having to store each of those data files locally.” Thus, *Sim* expressly attempts to avoid each network node from storing the full file locally, and as such one of ordinary skill in the art considering the disclosure of *Sim* as a whole, would not find sufficient reasoning to go against the express disclosure and goal of *Sim* to employ a distribution that results in each recipient node obtaining all of the subfiles (to result in a replicated volume of files such as is present in the system of *Zayas*).

On the other hand, *Zayas* is directed to a system in which full volumes of files are replicated on different servers and replication modules are employed for maintaining consistency among the volumes on the various servers. Thus, *Zayas* goes directly against the express teaching/desires of *Sim*. As such, insufficient reasoning exists for one of ordinary skill in the art to combine the teaching of *Zayas* with the teaching of *Sim* in the manner applied by the Examiner.

While the Examiner repeatedly acknowledges in the Final Office Action that there must be some reasoning why one skilled in the art would be motivated to make the proposed combination (see pages 2, 4, and 5 of the Final Office Action), the Examiner fails to actually identify any such reasoning. Instead, in response to the above arguments, the Final Office Action asserts on pages 4-5:

In this case, *Sim* discloses performing distribution of said file F to a plurality of recipient nodes using a distribution technique that comprises attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes [citations omitted]. *Zayas*, on the other hand, discloses an infrastructure for supporting file replications across a plurality of file servers for the purpose of automatically distributing data stored in files to users, thereby improving reliability, availability, local autonomy, load balancing, and data access performance.

While this asserts certain improvements that are allegedly provided by *Zayas*, such as improving reliability, availability, local autonomy, load balancing, and data access performance, this fails to explain any reasoning why one of ordinary skill in the art would go against the express goal of *Sim* of avoiding each network node from storing the full file locally in order to arrive at the asserted improvements offered by *Zayas*' system in which full volumes of data are stored at each node (and are updated to be maintained consistent with each other). Further, if *Zayas* system were desired to be employed (e.g., to arrive at the asserted improvements afforded by such a system in which each node has stored thereto a full volume of files), then one of ordinary skill in the art would be led away from employing the file distribution technique of *Sim* which expressly attempts to avoid each network node from storing the full file locally.

Thus, for at least this further reason, the rejections based on the applied combination of *Sim* and *Zayas* should be overturned.

Further, independent claim 17 recites, in part, “said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles” (emphasis added). However, *Chung* expressly teaches away from any such system that results in each recipient node attempting to obtain all of the subfiles. *Chung* expressly teaches at paragraphs 0005-0006 that replicating a full file onto a large number of servers is undesirable because it uses large amounts of expensive disk storage, etc. Thus, *Chung* proposes that a file be divided into a plurality of subfiles that are distributed to different servers without requiring that the entire file be distributed to each server.

On the other hand, *Zayas* is directed to a system in which full volumes of files are replicated on different servers and replication modules are employed for maintaining consistency among the volumes on the various servers. Thus, *Zayas* goes directly against the teaching of *Chung*. As such, one of ordinary skill in the art would not be motivated to combine the teaching of *Zayas* with the teaching of *Chung*. Thus, for this further reason, insufficient reasoning for one of ordinary skill in the art to combine *Chung* and *Zayas* exists, and thus the rejections that are based on a combination of *Chung* and *Zayas* should be overturned for this further reason.

Conclusion

In view of the above, Appellant requests that the board overturn the outstanding rejections of claims 1-28. Attached hereto are a Claims Appendix, Evidence Appendix, and Related Proceedings Appendix. As noted in the attached Evidence Appendix, no evidence pursuant to §§ 1.130, 1.131, or 1.132 or entered by or relied upon by the examiner is being submitted. Also, certain related appeals are identified in Section II above, but as noted by the Related Proceedings Appendix, no decisions have been received in such appeals and thus no copies of any decisions in related proceedings are provided.

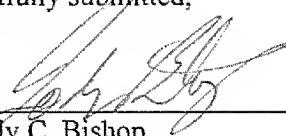
I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being transmitted via the Office electronic filing system in accordance with § 1.6(a)(4).

Dated: June 20, 2007

Signature: Donna Forbit
(Donna Forbit)

Respectfully submitted,

By:


Jody C. Bishop
Attorney/Agent for Applicant(s)
Reg. No. 44,034
Date: June 20, 2007
Telephone No. (214) 855-8007

VIII. CLAIMS APPENDIX

Claims Involved in the Appeal of Application Serial No. 10/619,737

1. A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

partitioning a file F into a plurality of subfiles;

performing distribution of said file F to a plurality of recipient nodes using a distribution technique that comprises

(a) attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of said first group but not all of said plurality of subfiles to any recipient node of said first group, and

(b) said plurality of recipient nodes of said first group attempting to exchange their respective subfiles received from said first node, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;

detecting a failed node of said plurality of recipient nodes; and

said distribution technique adapting to distribute all of the subfiles of said file F to each non-failed node of said plurality of recipient nodes.

2. The method of claim 1 wherein said distribution technique adapting responsive to said detecting a failed node.

3. The method of claim 1 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:

attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group.

4. The method of claim 1 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:

attempting to distribute the plurality of subfiles from said first node to said plurality of recipient nodes of said first group concurrently.

5. The method of claim 1 wherein said plurality of recipient nodes of said first group attempting to exchange their respective subfiles further comprises:

each of said plurality of recipient nodes attempting to establishing concurrent communication connections to every other recipient node of said first group.

6. The method of claim 1 wherein said detecting a failed node comprises said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node.

7. The method of claim 6 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises said first node attempting to establish concurrent communication connections to the recipient nodes of said first group, and wherein said distribution technique adapting comprises:

responsive to said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node, said first node using its established concurrent communication connections with non-failed nodes of said first group to communicate the particular subfile to said non-failed nodes.

8. The method of claim 6 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises said first node attempting to establish concurrent communication connections to the recipient nodes of said first group, and wherein said distribution technique adapting comprises:

responsive to said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node, said first node triggering a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate the particular subfile to said non-failed nodes.

9. The method of claim 1 wherein said detecting a failed node comprises:
said recipient nodes of said first group exchanging heartbeat messages;
at least one recipient node of said first group detecting a failed node from analysis of
heartbeat messages received; and
said at least one recipient node of said first group notifying said first node of said
detected failed node.

10. The method of claim 1 wherein said detecting a failed node comprises:
said non-failed recipient nodes of said first group sending heartbeat messages to said first
node; and
said first node detecting a failed node from analysis of received heartbeat messages from
said non-failed recipient nodes.

11. The method of claim 1 further comprising:
said first group of recipient nodes attempting to communicate said file *F* to a second
group comprising a plurality of recipient nodes.

12. The method of claim 11 further comprising:
each recipient node of said first group attempting to communicate a subfile to at least one
recipient node of said second group.

13. The method of claim 12 further comprising:
each recipient node of said first group attempting to communicate the subfile that it
received from said first node to a corresponding node of the second group.

14. The method of claim 12 wherein said detecting a failed node comprises detecting
a failed node in said first group when said failed node of said first group is attempting to
communicate a subfile to said at least one recipient node of said second group.

15. The method of claim 14 wherein said distribution technique adapting further comprises:

 said first node communicating said subfile to said at least one recipient node of said second group.

16. The method of claim 14 wherein said distribution technique adapting further comprises:

 said first node triggering a mirror node to communicate the subfile to said at least one recipient node of said second group.

17. A system comprising:

 an origin node operable to partition a file *F* into a plurality of subfiles, wherein said plurality of subfiles correspond in number to a number of recipient nodes in a first group to which said file is to be distributed;

 said origin node operable to attempt to distribute all of said plurality of subfiles to said recipient nodes, wherein said origin node attempts to distribute a different one of said plurality of subfiles to each of said recipient nodes;

 said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node;

 said origin node operable to detect a failed node in said first group; and

 said origin node operable to manage distribution of said file *F* upon detecting a failed node in said first group in a manner such that every non-failed node of said first group receives said file *F*.

18. The system of claim 17 wherein each of said recipient nodes are operable to attempt to distribute a subfile being received from said origin node to the others of said recipient nodes of said first group.

19. The system of claim 17 wherein said origin node is operable to attempt to distribute the plurality of subfiles to said plurality of recipient nodes of said first group concurrently.

20. The system of claim 17 wherein said origin node is operable to trigger a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate a subfile to said non-failed nodes.

21. A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

attempting to distribute a plurality of subfiles that comprise a file F from a first node to a first group comprising a plurality of recipient nodes, wherein the first node attempts to distribute at least one subfile to each recipient node of said first group but not all of said plurality of subfiles are distributed from the first node to any of the recipient nodes of said first group;

said plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;

detecting whether one of said plurality of recipient nodes of said first group has failed; and

if a recipient node of said first group has failed, managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that the file F is distributed to each non-failed node of said plurality of recipient nodes.

22. The method of claim 21 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:

attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group.

23. The method of claim 22 wherein said managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that file *F* is distributed to each non-failed node of said plurality of recipient nodes comprises:

 said first node communicating to non-failed nodes of said first group a subfile that the first node would communicate to the failed node if the failed node were not failed.

24. The method of claim 22 wherein said managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that file *F* is distributed to each non-failed node of said plurality of recipient nodes comprises:

 said first node triggering a mirror node to communicate to non-failed nodes of said first group a subfile that the first node would communicate to the failed node if the failed node were not failed.

25. The method of claim 21 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:

 attempting to distribute the plurality of subfiles from said first node to said plurality of recipient nodes of said first group concurrently.

26. The method of claim 21 wherein said plurality of recipient nodes of said first group attempting to exchange their respective subfiles further comprises:

 each of said plurality of recipient nodes attempting to establishing concurrent communication connections to every other recipient node of said first group.

27. The method of claim 21 wherein said detecting whether one of said plurality of recipient nodes of said first group has failed comprises:

 said recipient nodes of said first group exchanging heartbeat messages;

 at least one recipient node of said first group detecting a failed node from analysis of heartbeat messages received; and

 said at least one recipient node of said first group notifying said first node of said detected failed node.

28. The method of claim 21 wherein said detecting whether one of said plurality of recipient nodes of said first group has failed comprises:

the non-failed recipient nodes of said first group sending heartbeat messages to said first node; and

 said first node detecting a failed node from analysis of received heartbeat messages from the non-failed recipient nodes.

IX. EVIDENCE APPENDIX

No evidence pursuant to §§ 1.130, 1.131, or 1.132 or entered by or relied upon by the examiner is being submitted.

X. RELATED PROCEEDINGS APPENDIX

Appellant respectfully notes that the following copending applications are on appeal before the Board, which contain at least some issues that are similar to issues of the present application, which may be affected or have a bearing on the Board's decision in this appeal: 1) Application No. 10/619,805 (hereinafter "the '805 application"); 2) Application No. 10/345,716 (hereinafter "the '716 application"); 3) Application No. 10/345,587 (hereinafter "the '587 application"); 4) Application No. 10/345,718 (hereinafter "the '718 application"); and 5) Application No. 10/345,719 (hereinafter "the '719 application").

As of the filing of this Appeal Brief, no decisions have been received in these related appeals, and thus no copies of such decisions in the related proceedings are provided.